

Graph spectra and the detectability of community structure in networks

Raj Rao Nadakuditi¹ and M. E. J. Newman²

¹*Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI 48109*

²*Department of Physics and Center for the Study of Complex Systems, University of Michigan, Ann Arbor, MI 48109*

We study networks that display community structure—groups of nodes within which connections are unusually dense. Using methods from random matrix theory, we calculate the spectra of such networks in the limit of large size, and hence demonstrate the presence of a phase transition in matrix methods for community detection, such as the popular modularity maximization method. The transition separates a regime in which such methods successfully detect the community structure from one in which the structure is present but is not detected. By comparing these results with recent analyses of maximum-likelihood methods we are able to show that spectral modularity maximization is an optimal detection method in the sense that no other method will succeed in the regime where the modularity method fails.

The problem of community detection in networks has attracted a substantial amount of attention in recent years [1, 2]. Communities in this context are groups of vertices within a network that have a high density of within-group connections but a lower density of between-group connections. The challenge is to find such groups accurately and efficiently in a given network—the ability to do so would have applications in the analysis of observational data, network visualization, and complexity reduction and parallelization of network problems.

In this paper we focus on matrix methods for community detection, which are based on the properties of matrix representations of networks such as the adjacency matrix or the modularity matrix. While significant effort has been devoted to the development of practical algorithms using these methods, there has been less work on formal examination of their properties and implications for algorithm performance. Here we give an analysis of the spectral properties of the adjacency and modularity matrices using random matrix methods, and in the process uncover a number of results of practical importance. Chief among these is the presence of a sharp transition between a regime in which the spectrum contains clear evidence of community structure and a regime in which it contains none. In the former regime, community detection is possible and current algorithms should perform well; in the latter, any method relying on the spectrum to perform structure detection must fail. A similar phase transition has been reported recently in an analysis of a different class of detection methods, based on Bayesian inference [3]. By comparing the two analyses, we are able to demonstrate that methods such as modularity maximization are optimal, in the sense that no other method will succeed where they fail.

For the formal analysis of community structured networks, we must define the particular network or networks we will study. In this paper we focus on the most widely studied model of community structure, the stochastic block model, although our methods could be applied to other models as well. The stochastic block model, in its simplest form, divides a network of n vertices into

some number q of groups denoted by $r = 1 \dots q$ and then places undirected edges between vertex pairs i, j with independent probabilities p_{rs} , where r, s are respectively the groups to which vertices i, j belong. In other words, the probability of an edge between two vertices in this model depends only on the groups in which the vertices fall. If the diagonal elements of the matrix of probabilities p_{rs} are greater than the off-diagonal elements, then the network displays classic community structure with a greater density of edges within groups than between them. Particular instances of the stochastic block model are commonly used as testbeds for assessing the performance of community detection algorithms—especially in the “four groups” test [1] and the planted partition model [4].

Let us first demonstrate our argument for the simplest possible case of a network with $q = 2$ groups of equal size $\frac{1}{2}n$ each and just two different probabilities p_{in} and p_{out} for connections within and between groups. We focus particularly on the case of sparse networks, those for which the fraction of possible edges that are present in the network vanishes in the limit of large n , which appears to be representative of most networks observed in the real world, although our results apply in principle to dense networks as well.

The adjacency matrix \mathbf{A} of an undirected network is the $n \times n$ symmetric matrix with elements $A_{ij} = 1$ if vertices i and j are connected by an edge and 0 otherwise. If we average the adjacency matrix over the ensemble of our stochastic block model the resulting matrix $\langle \mathbf{A} \rangle$ has elements equal to p_{in} for vertices in the same group and p_{out} for vertices in different groups. Defining $c_{\text{in}} = np_{\text{in}}$ and $c_{\text{out}} = np_{\text{out}}$, this matrix can be written in the form

$$\langle \mathbf{A} \rangle = \frac{1}{2}(c_{\text{in}} + c_{\text{out}}) \mathbf{1}\mathbf{1}^T + \frac{1}{2}(c_{\text{in}} - c_{\text{out}}) \mathbf{u}\mathbf{u}^T, \quad (1)$$

where $\mathbf{1}$ and \mathbf{u} are the unit vectors $\mathbf{1} = (1, 1, 1, \dots)/\sqrt{n}$ and $\mathbf{u} = (1, 1, \dots, -1, -1, \dots)/\sqrt{n}$, the ± 1 elements in the latter denoting the members of the two communities.

Now the full adjacency matrix can be written in the form $\mathbf{A} = \langle \mathbf{A} \rangle + \mathbf{X}$ where the matrix \mathbf{X} is the deviation

between the adjacency matrix and its average value. By definition, \mathbf{X} is a symmetric random matrix with independent elements of mean zero.

Our analysis will focus on the spectrum of eigenvalues z of the adjacency matrix, which we calculate in several steps. We start by calculating the spectral density $\rho(z)$ of the matrix \mathbf{X} alone, whose average value in the random ensemble can be written in terms of the imaginary part of the Stieltjes transform:

$$\rho(z) = -\frac{1}{\pi} \text{Im} \langle \text{Tr} (z\mathbf{I} - \mathbf{X})^{-1} \rangle, \quad (2)$$

where $\langle \dots \rangle$ indicates the ensemble average. The average of the trace can be expanded in powers of \mathbf{X} as

$$\langle \text{Tr}(z\mathbf{I} - \mathbf{X})^{-1} \rangle = \frac{1}{z} \sum_{k=0}^{\infty} \frac{\text{Tr} \langle \mathbf{X}^k \rangle}{z^k}, \quad (3)$$

where the individual terms take the form

$$\text{Tr} \langle \mathbf{X}^k \rangle = \sum_{i_1 \dots i_k} \langle X_{i_1 i_2} X_{i_2 i_3} \dots X_{i_k i_1} \rangle. \quad (4)$$

Since the elements of \mathbf{X} have mean zero, any term in this sum that contains any variable just once will average to zero. Moreover, terms containing any variable more than twice become negligible when the average degree of the network is much greater than one, so that the only terms remaining are those for which k is even and which contain each variable exactly twice. Geometrically, the sequence of indices in these terms takes the form of an Euler tour of a rooted plane tree, with a factor of $\langle X_{ij}^2 \rangle$ on each edge, whose average value is $\frac{1}{2}(p_{\text{in}} + p_{\text{out}}) = (c_{\text{in}} + c_{\text{out}})/2n$. Writing $k = 2m$ with m integer, there are n^{m+1} ways to choose the $m+1$ vertices of the tree and the number of topologically distinct rooted plane trees with this many vertices is equal to the Catalan number C_m . Thus

$$\begin{aligned} \text{Tr} \langle \mathbf{X}^{2m} \rangle &= n^{m+1} \left(\frac{c_{\text{in}} + c_{\text{out}}}{2n} \right)^m C_m \\ &= n \left[\frac{1}{2}(c_{\text{in}} + c_{\text{out}}) \right]^m C_m. \end{aligned} \quad (5)$$

Combining this result with Eq. (3), we have

$$\begin{aligned} \langle \text{Tr}(z\mathbf{I} - \mathbf{X})^{-1} \rangle &= \frac{n}{z} \sum_{m=0}^{\infty} \left[\frac{1}{2}(c_{\text{in}} + c_{\text{out}}) \right]^m C_m / z^{2m} \\ &= \frac{n}{c_{\text{in}} + c_{\text{out}}} \left[z - \sqrt{z^2 - 2(c_{\text{in}} + c_{\text{out}})} \right]. \end{aligned} \quad (6)$$

Then the spectral density, Eq. (2), is

$$\rho(z) = (n/\pi) \frac{\sqrt{2(c_{\text{in}} + c_{\text{out}}) - z^2}}{c_{\text{in}} + c_{\text{out}}}, \quad (7)$$

which is a modified form of the classic Wigner semicircle law for random matrices. Note that the density of

eigenvalues increases with n , which implies that the fluctuations in the values vanish as $n \rightarrow \infty$.

Armed with this result we can now calculate the spectrum of the adjacency matrix $\mathbf{A} = \langle \mathbf{A} \rangle + \mathbf{X}$, but again we take the calculation in stages, starting with the simpler exercise of calculating the spectrum of the matrix

$$\mathbf{B} = \frac{1}{2}(c_{\text{in}} - c_{\text{out}}) \mathbf{u}\mathbf{u}^T + \mathbf{X} = \mathbf{A} - \frac{1}{2}(c_{\text{in}} + c_{\text{out}}) \mathbf{1}\mathbf{1}^T. \quad (8)$$

Note that $\frac{1}{2}(c_{\text{in}} + c_{\text{out}}) \mathbf{1}\mathbf{1}^T$ is the uniform matrix with all elements equal to $\frac{1}{2}(p_{\text{in}} + p_{\text{out}})$, which is the average probability p of an edge in the entire network. Hence the elements of \mathbf{B} are $B_{ij} = A_{ij} - p$. This matrix is of interest in its own right. It is the so-called modularity matrix, which forms the basis for the modularity maximization method of community detection. The modularity matrix is usually defined by $B_{ij} = A_{ij} - P_{ij}$ where P_{ij} is the expected value of the adjacency matrix element in a null model containing no community structure. The most commonly used null model is the configuration model, a random graph with specified degree distribution, but in the present case, for which all vertices have the same expected degree, the null model is just a standard Erdős–Rényi random graph with $P_{ij} = p$ for all i, j , leading to the definition in Eq. (8). Thus our calculation will in this case give us also the spectrum of the modularity matrix.

The general form of the matrix \mathbf{B} is that of a rank-1 matrix $\mathbf{u}\mathbf{u}^T$ plus a random perturbation, a form that has been studied in the mathematical literature. Following an argument of [5, 6], let z be an eigenvalue of this matrix and \mathbf{v} be the corresponding normalized eigenvector, so that

$$\left[\frac{1}{2}(c_{\text{in}} - c_{\text{out}}) \mathbf{u}\mathbf{u}^T + \mathbf{X} \right] \mathbf{v} = z\mathbf{v}. \quad (9)$$

A rearrangement gives $(z\mathbf{I} - \mathbf{X})\mathbf{v} = \frac{1}{2}(c_{\text{in}} - c_{\text{out}}) \mathbf{u}\mathbf{u}^T \mathbf{v}$, where \mathbf{I} is the identity. Multiplying by $\mathbf{u}^T(z\mathbf{I} - \mathbf{X})^{-1}$ and cancelling a factor of $\mathbf{u}^T \mathbf{v}$, we find that

$$\frac{2}{c_{\text{in}} - c_{\text{out}}} = \mathbf{u}^T (z\mathbf{I} - \mathbf{X})^{-1} \mathbf{u} = \sum_{i=1}^n \frac{(\mathbf{u}^T \mathbf{x}_i)^2}{z - \lambda_i}, \quad (10)$$

where λ_i is the i th eigenvalue of \mathbf{X} and \mathbf{x}_i is the corresponding eigenvector.

The solutions of this equation, which give the eigenvalues z of the modularity matrix, are represented graphically in Fig. 1a. The right-hand side of the equation has poles at $z = \lambda_i$ for all i and, as the figure shows, this means that the eigenvalues must satisfy $z_1 \geq \lambda_1 \geq z_2 \geq \lambda_2 \geq \dots \geq z_n \geq \lambda_n$, where both sets of eigenvalues are numbered in order from largest to smallest. These inequalities place bounds on the eigenvalues $z_2 \dots z_n$ that become tight as $n \rightarrow \infty$, meaning that the spectrum of the modularity matrix is asymptotically identical to that of the random matrix \mathbf{X} .

The only exception is the highest eigenvalue z_1 , which is bounded below by λ_1 but unbounded above. To calculate this eigenvalue we note that since \mathbf{X} is a random matrix, its eigenvectors are also random, so that cross-terms

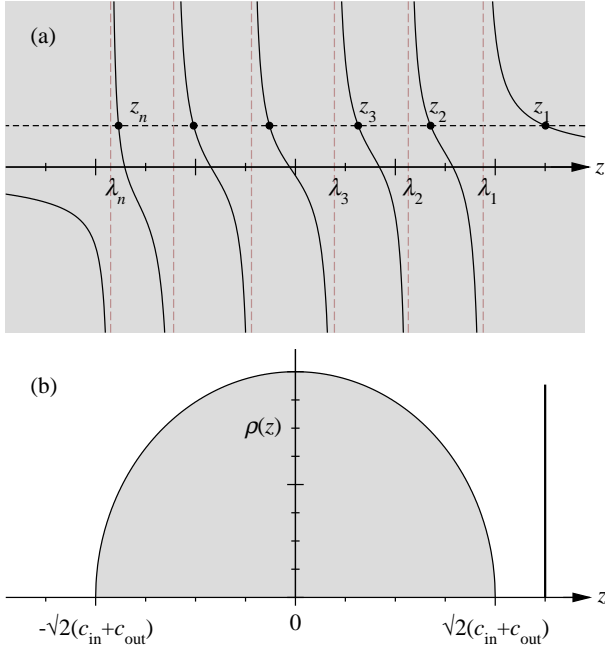


FIG. 1: (a) The solid curve represents the right-hand side of Eq. (10) while the dashed horizontal line represents the left-hand side. The points at which the two cross, indicated by the dots, are the solutions z_i of the equation, which necessarily fall between the eigenvalues λ_i of the matrix \mathbf{X} (vertical dashed lines). (b) The spectrum of the modularity matrix is the same of that of the random matrix \mathbf{X} (semicircle), except for the highest eigenvalue z_1 , which is separate from the rest of the spectrum above the transition point given in Eq. (14).

cancel in the quantity $(\mathbf{u}^T \mathbf{x}_i)^2$ and the average value is simply $|\mathbf{x}_i|^2/n = 1/n$. Taking the average of (10) over the random matrix ensemble in the limit of large n then gives

$$\begin{aligned} \frac{2}{c_{\text{in}} - c_{\text{out}}} &= \frac{1}{n} \left\langle \sum_{i=1}^n \frac{1}{z - \lambda_i} \right\rangle = \frac{1}{n} \langle \text{Tr}(z\mathbf{I} - \mathbf{X})^{-1} \rangle \\ &= \frac{z - \sqrt{z^2 - 2(c_{\text{in}} + c_{\text{out}})}}{c_{\text{in}} + c_{\text{out}}}, \end{aligned} \quad (11)$$

where we have used Eq. (6). Rearranging for z , we get our expression for the leading eigenvalue z_1 :

$$z_1 = \frac{1}{2}(c_{\text{in}} - c_{\text{out}}) + \frac{c_{\text{in}} + c_{\text{out}}}{c_{\text{in}} - c_{\text{out}}}. \quad (12)$$

We can use the same method to deduce the spectrum of the full adjacency matrix also. From Eq. (8) we see that the adjacency matrix takes the form $\mathbf{A} = \frac{1}{2}(c_{\text{in}} + c_{\text{out}})\mathbf{1}\mathbf{1}^T + \mathbf{B}$, which is again a rank-1 matrix plus a random perturbation. By the same argument as before, we can show that this matrix has all eigenvalues the same (to within tight bounds) as those of the modularity matrix, except again for the leading eigenvalue, whose value can be calculated from a relation of

the form (10). The end result is that the lower $n - 2$ eigenvalues of the adjacency matrix have the same spectrum as the random matrix \mathbf{X} and the top two have the values z_1 , Eq. (12), and

$$z_2 = \frac{1}{2}(c_{\text{in}} + c_{\text{out}}) + 1. \quad (13)$$

With this result, we now have the complete spectrum for both the adjacency matrix and the modularity matrix.

Let us focus on the modularity matrix. The spectrum is depicted in Fig. 1b and consists of the continuous semicircular band of eigenvalues, Eq. (7), plus the single eigenvalue z_1 , Eq. (12). If the network contained no community structure, then z_1 would not be separated from the continuous band as it is here. So long as it is well separated the spectrum shows clear evidence of the existence of community structure and one can reasonably say that a calculation of the spectrum constitutes positive “detection” of that structure. Moreover, the signs of the elements of the leading eigenvector provide a good guide to the community division of the network, and indeed this particular method for community identification can be derived directly as a spectral version of the standard method of modularity maximization [7]. If, however, the position of the leading eigenvalue passes the edge of the continuous band, the spectrum no longer shows evidence of community structure and spectral algorithms based on the corresponding eigenvector will fail. One might imagine that this point would arrive when $c_{\text{in}} = c_{\text{out}}$, which is the point at which the network contains no community structure at all, but this is not the case. From Eq. (7) we see that the end of the continuous band falls at $z = \sqrt{2(c_{\text{in}} + c_{\text{out}})}$ and, setting z_1 from Eq. (12) equal to this value, we find that we lose the ability to detect community structure at an earlier point, when

$$c_{\text{in}} - c_{\text{out}} = \sqrt{2(c_{\text{in}} + c_{\text{out}})}. \quad (14)$$

This value sets a detectability threshold beyond which the communities are present but cannot be detected. For $c_{\text{in}} - c_{\text{out}}$ smaller than this value, but greater than zero, community structure is present in the network in the sense that the average probability of edges within groups is measurably higher than that between groups, but we nonetheless fail to find the communities using our spectral method. One can generalize the calculation to networks with a larger number q of communities and we find that a similar transition happens at the point

$$c_{\text{in}} - c_{\text{out}} = \sqrt{q[c_{\text{in}} + (q-1)c_{\text{out}}]}. \quad (15)$$

The existence of a transition of this kind, though not its precise location, was demonstrated previously using different methods by Reichardt and Leone [8] and there are also close connections between our calculation and the theory of disordered systems [9].

One might imagine this transition to be a particular property of the spectral method we have considered. Perhaps a different modularity maximization algorithm, one

not based on spectral techniques, or a different type of community detection method altogether, would be able to get past this detectability threshold. This, however, is also not the case.

In recent work, Decelle *et al.* [3] have used arguments based on the cavity method of statistical physics to demonstrate the existence of a transition akin to the one above in another community detection method, a Bayesian maximum-likelihood method based on directly fitting the stochastic block model to a network. Moreover, their transition falls at the same position as that of Eq. (15). The importance of this result stems from the fact that if we know the model from which a network is drawn, then fitting directly to that model is provably the optimal way of recovering the parameters of the model used to generate the network—including, in this case, the community structure. Thus, as Decelle *et al.* have pointed out, their maximum-likelihood method is an optimal method in the sense that no method can detect communities in the regime where their method fails. Unfortunately fitting to the stochastic block model turns out to be a poor method of community detection for real-world networks [3, 10], but its optimality in the present case is a useful result nonetheless. It implies, given that the detectability transition falls in the same place as for the spectral modularity method, that the modularity method is also optimal in the same sense: no other method will detect communities in the network when the modularity method does not [11].

We can take these calculations further. For instance, we can calculate the expected fraction of vertices classified correctly by the spectral algorithm. We can show that the elements of the leading eigenvector \mathbf{v} of the modularity matrix are equal to $\pm\alpha/\sqrt{n}$ plus Gaussian perturbations with variance $(1 - \alpha^2)/n$, where

$$\alpha^2 = \frac{(c_{\text{in}} - c_{\text{out}})^2 - 2(c_{\text{in}} + c_{\text{out}})}{(c_{\text{in}} - c_{\text{out}})^2}. \quad (16)$$

Then the fraction of elements that retain the correct sign and hence give correct classifications of the corresponding vertices is $\frac{1}{2}[1 + \text{erf} \sqrt{\alpha^2/2(1 - \alpha^2)}]$, where $\text{erf} x$ is the Gaussian error function.

Figure 2 shows a plot of this quantity as a function of $c_{\text{in}} - c_{\text{out}}$ for networks with several different values of the average degree, along with results for the same quantity from actual applications of the spectral modularity algorithm to networks generated using the stochastic block model. As the figure shows, the agreement between the two is excellent, except in the immediate vicinity of the phase transition, where finite-size effects produce some rounding of the threshold. The fraction of correctly classified vertices (minus $\frac{1}{2}$) plays the role of an order parameter for the detectability transition. Since it is continuous at the transition point, we have a continuous phase transition.

The calculations presented here could be extended in




FIG. 2: The fraction of vertices correctly classified by the spectral modularity algorithm in networks generated using the block model studied here, as a function of $c_{\text{in}} - c_{\text{out}}$, for four different values of the average degree as indicated. Points are numerical measurements for networks of 100 000 vertices, averaged over 25 networks each; solid curves represent the analytic prediction. The phase transition at which the algorithm fails is clearly visible in each curve.

a number of additional directions. For instance, the results given are accurate for networks with large average degree but for networks with smaller degree there are additional corrections that corresponding to additional terms in the trace, Eq. (5). A calculation of these sub-leading terms would help to complete the picture for low-degree networks. Also our calculations all use the standard stochastic block model, and although this is the model most widely used for benchmark calculations and synthetic tests, other models have been proposed, such as the degree-corrected block model [10] or more exotic models such as the LFR benchmark networks [12]. It would be useful to know if results similar to those described here can be derived for these more complex models.

The authors thank Cris Moore and Lenka Zdeborova for useful comments. This work was funded in part by the National Science Foundation under grants CCF-1116115 and DMS-1107796 and by the Office of Naval Research under grant N00014-11-1-0660.

-
- [1] M. Girvan and M. E. J. Newman, Community structure in social and biological networks. *Proc. Natl. Acad. Sci. USA* **99**, 7821–7826 (2002).
 - [2] S. Fortunato, Community detection in graphs. *Phys. Rep.* **486**, 75–174 (2010).
 - [3] A. Decelle, F. Krzakala, C. Moore, and L. Zdeborová, Inference and phase transitions in the detection of modules in sparse networks. *Phys. Rev. Lett.* **107**, 065701 (2011).

- [4] A. Condon and R. M. Karp, Algorithms for graph partitioning on the planted partition model. *Random Structures and Algorithms* **18**, 116–140 (2001).
- [5] F. Benaych-Georges and R. R. Nadakuditi, The eigenvalues and eigenvectors of finite, low rank perturbations of large random matrices. *Advances in Mathematics* **227**, 494–521 (2011).
- [6] M. Capitaine, C. Donati-Martin, and D. Féral, The largest eigenvalues of finite rank deformation of large Wigner matrices: Convergence and nonuniversality of the fluctuations. *Annals of Probability* **37**, 1–47 (2009).
- [7] M. E. J. Newman, Modularity and community structure in networks. *Proc. Natl. Acad. Sci. USA* **103**, 8577–8582 (2006).
- [8] J. Reichardt and M. Leone, (Un)detectable cluster structure in sparse networks. *Phys. Rev. Lett.* **101**, 078701 (2008).
- [9] J. M. Kosterlitz, D. J. Thouless, and R. C. Jones, Spherical model of a spin-glass. *Phys. Rev. Lett.* **36**, 1217–1220 (1976).
- [10] B. Karrer and M. E. J. Newman, Stochastic blockmodels and community structure in networks. *Phys. Rev. E* **83**, 016107 (2011).
- [11] Strictly speaking, the results of Ref. [3] say only that no algorithm will succeed that works in polynomial time—it may be possible, at least in certain parameter regimes, for non-polynomial algorithms to detect the community structure. Non-polynomial algorithms, however, are prohibitively slow for all but the smallest of networks, making communities effectively undetectable past the transition point.
- [12] A. Lancichinetti, S. Fortunato, and F. Radicchi, Benchmark graphs for testing community detection algorithms. *Phys. Rev. E* **78**, 046110 (2008).